

Exhibit 8 to Complaint
Intellectual Ventures I LLC and Intellectual Ventures II LLC

Example Southwest Count 1 Systems and Services
U.S. Patent No. 8,332,844 (“the ’844 Patent”)

The Accused Systems and Services include without limitation Southwest systems and services that utilize Docker; all past, current, and future systems and services that operate in the same or substantially similar manner as the specifically identified systems and services; and all past, current, and future Southwest systems and services that have the same or substantially similar features as the specifically identified systems and services (“Example Southwest Count 1 Systems and Services” or “Southwest Systems and Services”).¹

On information and belief, the Southwest Systems and Services use Docker in its private cloud(s). For example, Southwest posts, or has posted, job opportunities that require familiarity with Docker containerization concepts.

See <https://www.linkedin.com/in/charlesmarshall-eth/>, the job profile of an associate software engineer listing Docker as a skill. (last accessed 9/23/2024).

See <https://www.linkedin.com/in/bhaveskhar-kongari/>, the job profile of a senior DevOps Engineer listing usage of Docker. (last accessed 9/24/24).

As another example, Southwest has stated that it is investing in cloud technology and has “moved about 50% of its technology” to the cloud and has indicated cloud migration is one of its areas of focus for 2024 and beyond. Source: <https://www.phocuswire.com/southwest-airlines-cio-tech-investment>.

¹ For the avoidance of doubt, Plaintiffs do not accuse public clouds of Southwest if those services are provided by a cloud provider with a license to Plaintiffs’ patents that covers Southwest’s activities. IV will provide relevant license agreements for cloud providers in discovery. To the extent any of these licenses are relevant to Southwest’s activities, Plaintiffs will meet and confer with Southwest about the impact of such license(s).

On information and belief, other information confirms Southwest uses Docker technology.



Top Airlines, Airports & Air Services Companies Using Docker

37,841 companies using this technology

By [Docker](#)

Docker is a software container platform. Developers use Docker to eliminate “works on my machine” problems when collaborating on code with co-workers. Operators use Docker to run and manage apps side-by-side in isolated containers to get better compute density. Enterprises use Docker to build agile software delivery pipelines. [Read less](#)



Southwest Airlines

Technologies used by the company: 1,161

Source: <https://www.zoominfo.com/tech/23717/docker-tech-from-transportation-airline-industry-in-us-by-revenue>.²

² All sources cited in this document were publicly accessible as of the filing date of the Complaint.

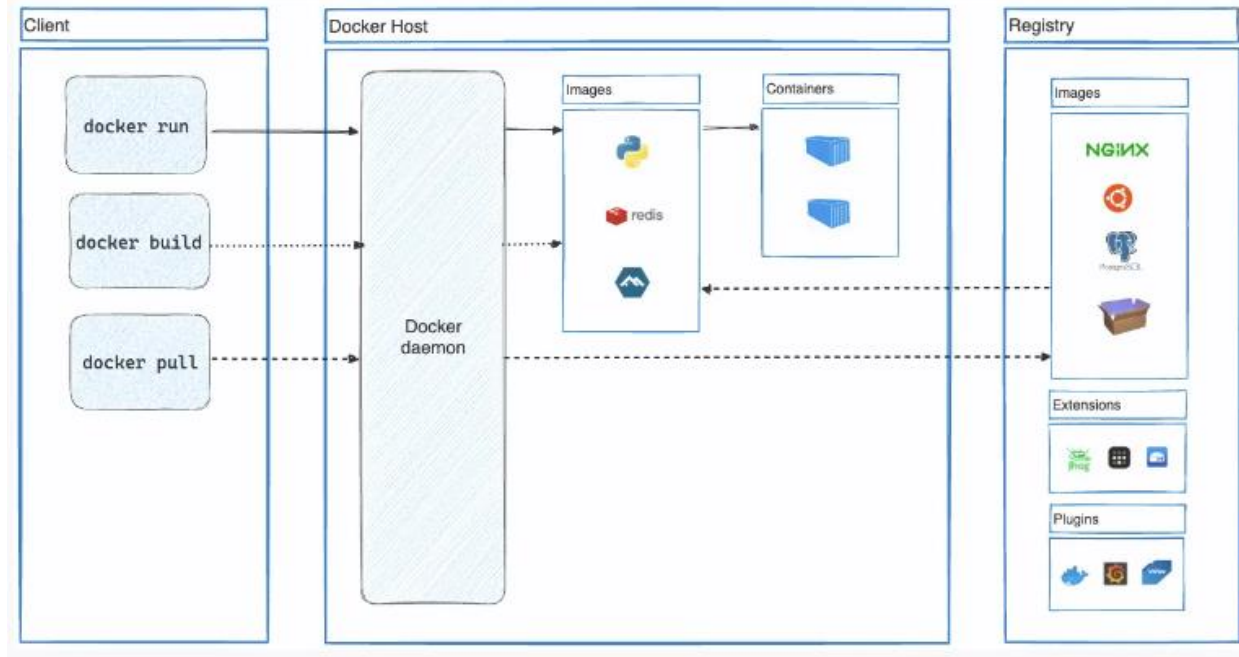
U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
[7.pre]. A method for providing data to a plurality of compute nodes, comprising:	<p>To the extent this preamble is limiting, on information and belief, the Southwest Count 1 Systems and Services practice a method for providing data to a plurality of compute nodes.</p> <p>For example, on information and belief, Southwest's use of Docker provides "a method for providing data to a plurality of compute nodes."</p> <p>Docker is an open platform for developing, shipping, and running applications. Docker enables you to separate your applications from your infrastructure so you can deliver software quickly. With Docker, you can manage your infrastructure in the same ways you manage your applications. By taking advantage of Docker's methodologies for shipping, testing, and deploying code, you can significantly reduce the delay between writing code and running it in production.</p> <p>Source: https://docs.docker.com/get-started/overview/.</p> <h2>The Docker platform</h2> <p>Docker provides the ability to package and run an application in a loosely isolated environment called a container. The isolation and security lets you run many containers simultaneously on a given host. Containers are lightweight and contain everything needed to run the application, so you don't need to rely on what's installed on the host. You can share containers while you work, and be sure that everyone you share with gets the same container that works in the same way.</p> <p>Source: https://docs.docker.com/get-started/overview/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Docker provides tooling and a platform to manage the lifecycle of your containers:</p> <ul style="list-style-type: none"> • Develop your application and its supporting components using containers. • The container becomes the unit for distributing and testing your application. • When you're ready, deploy your application into your production environment, as a container or an orchestrated service. This works the same whether your production environment is a local data center, a cloud provider, or a hybrid of the two. <p>Source: https://docs.docker.com/get-started/overview/.</p> <p>On information and belief, Docker uses a client-server architecture. For example, the Docker architecture shown below shows a client, a Docker host, and a registry. Docker clients communicate with a Docker daemon, which builds, runs, and distributes Docker containers. Docker images are run within containers.</p>

U.S. Patent No. 8,332,844 (Claim 7)

Claim(s)

Example Southwest Count 1 Systems and Services



Source: <https://docs.docker.com/get-started/overview/>.

[7.a] storing blocks of a root image of said compute nodes on a first storage unit;

On information and belief, the Southwest Count 1 Systems and Services practice storing blocks of a root image of said compute nodes on a first storage unit.

For example, on information and belief, Docker stores blocks of a root image including a collection of read-only layers in a Docker image in the Docker registry.


U.S. Patent No. 8,332,844 (Claim 7)

Claim(s)

Example Southwest Count 1 Systems and Services

The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an `ubuntu:15.04` image.

Source: <https://docs.docker.com/storage/storagedriver/>.




Source: <https://docs.docker.com/storage/storagedriver/>.

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Distribution Registry</p> <p>What it is</p> <p>The Registry is a stateless, highly scalable server side application that stores and lets you distribute container images and other content. The Registry is open-source, under the permissive Apache license.</p> <p>Source: https://docs.docker.com/registry/.</p> <p>Why use it</p> <p>You should use the Registry if you want to:</p> <ul style="list-style-type: none"> • tightly control where your images are being stored • fully own your images distribution pipeline • integrate image storage and distribution tightly into your in-house development workflow <p>Source: https://docs.docker.com/registry/.</p> <p>On information and belief, the image below shows blocks of a root image (associated with ubuntu version 15.04) that is associated with a compute node.</p>

U.S. Patent No. 8,332,844 (Claim 7)

Claim(s)

Example Southwest Count 1 Systems and Services



Source: <https://docs.docker.com/storage/storagedriver/>.

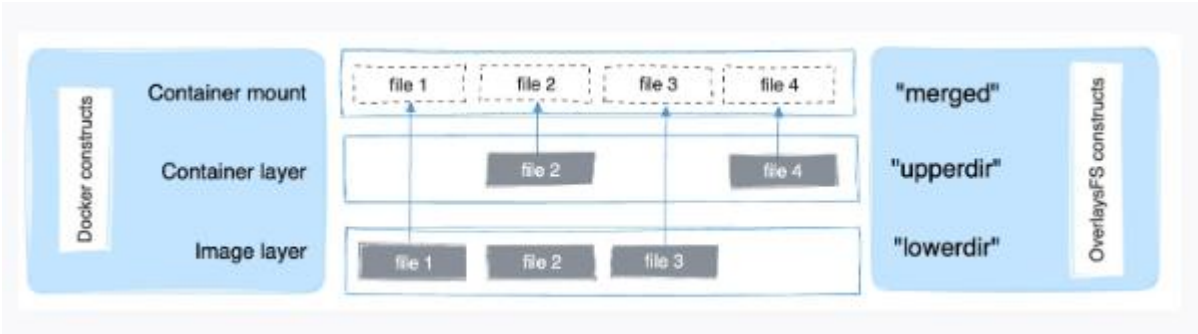
On information and belief, an image stores blocks of a root image, such as an ubuntu image or other operating systems software and other applications and software.

An image is a read-only template with instructions for creating a Docker container. Often, an image is based on another image, with some additional customization. For example, you may build an image which is based on the `ubuntu` image, but installs the Apache web server and your application, as well as the configuration details needed to make your application run.

Source: <https://docs.docker.com/get-started/overview/>.

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Docker registries</p> <p>A Docker registry stores Docker images. Docker Hub is a public registry that anyone can use, and Docker looks for images on Docker Hub by default. You can even run your own private registry.</p> <p>Source: https://docs.docker.com/get-started/overview/.</p> <p>On information and belief, Docker includes images that are built up from a series of layers, where each layer represents an instruction in the image's Dockerfile. Blocks of a root image are stored in a memory associated with a computer such as a server (node).</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<h2>Images and layers</h2> <p>A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p>This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted, the writable layer is also deleted. The underlying image remains unchanged.</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>On information and belief, a Docker image and a Docker container can be layered in a manner as shown in the image below.</p>  <p>Source: https://docs.docker.com/engine/storage/drivers/overlayfs-driver/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<i>See also</i> preamble [7.pre].
[7.b] storing leaf images for respective compute nodes on respective second storage units,	<p>On information and belief, the Southwest Count 1 Systems and Services practice storing leaf images for respective compute nodes on respective second storage units.</p> <p>For example, on information and belief, Docker stores leaf images including thin R/W layer in a runnable container for respective compute nodes on respective second storage units including Docker's local storage area.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an <code>ubuntu:15.04</code> image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>When you use <code>docker pull</code> to pull down an image from a repository, or when you create a container from an image that doesn't yet exist locally, each layer is pulled down separately, and stored in Docker's local storage area, which is usually <code>/var/lib/docker/</code> on Linux hosts. You can see these layers being pulled in this example:</p> <pre> \$ docker pull ubuntu:22.04 22.04: Pulling from library/ubuntu f476d66f5408: Pull complete 8882c27f669e: Pull complete d9af21273955: Pull complete f5029279ec12: Pull complete Digest: sha256:6120be6a2b7ce665d0cbddc3ce6eae60fe94637c6a66985312d11 Status: Downloaded newer image for ubuntu:22.04 docker.io/library/ubuntu:22.04 </pre> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>


U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<h2>Images and layers</h2> <p>A Docker image is built up from a series of layers. Each layer represents an instruction in the image's Dockerfile. Each layer except the very last one is read-only. Consider the following Dockerfile:</p> <pre># syntax=docker/dockerfile:1 FROM ubuntu:22.04 LABEL org.opencontainers.image.authors="org@example.com" COPY . /app RUN make /app RUN rm -r \$HOME/.cache CMD python /app/app.py</pre> <p>This Dockerfile contains four commands. Commands that modify the filesystem create a layer. The <code>FROM</code> statement starts out by creating a layer from the <code>ubuntu:22.04</code> image. The <code>LABEL</code> command only modifies the image's metadata, and doesn't produce a new layer. The <code>COPY</code> command adds some files from your Docker client's current directory. The first <code>RUN</code> command builds your application using the <code>make</code> command, and writes the result to a new layer. The second <code>RUN</code> command removes a cache directory, and writes the result to a new layer. Finally, the <code>CMD</code> instruction specifies what command to run within the container, which only modifies the image's metadata, which doesn't produce an image layer.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted, the writable layer is also deleted. The underlying image remains unchanged.</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)

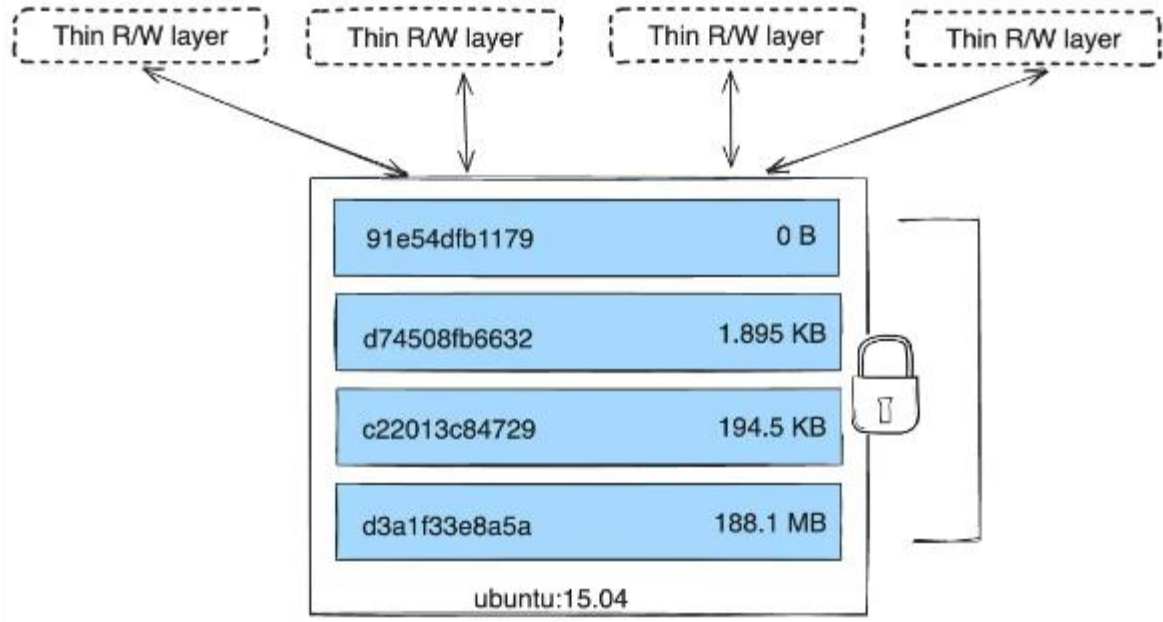
Claim(s)

Example Southwest Count 1 Systems and Services

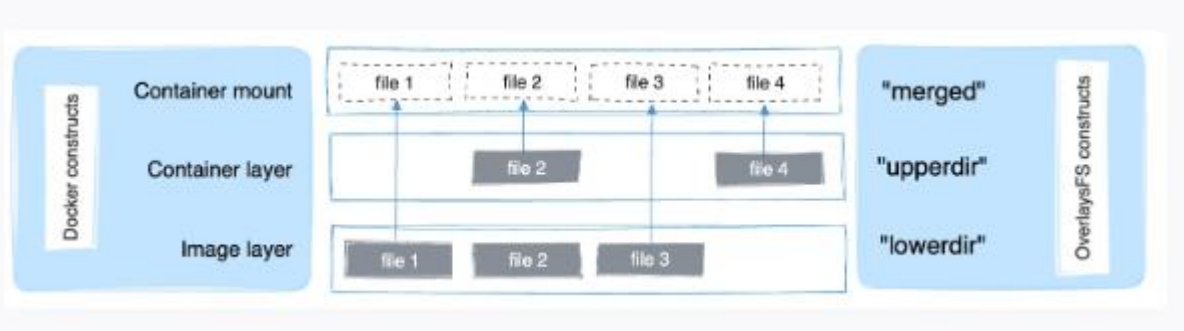


Source: <https://docs.docker.com/engine/storage/drivers/>.

On information and belief, the image below shows a thin r/w layer for each of four separate containers that are each stored in memory and are associated with a root image that is further associated with a compute node.

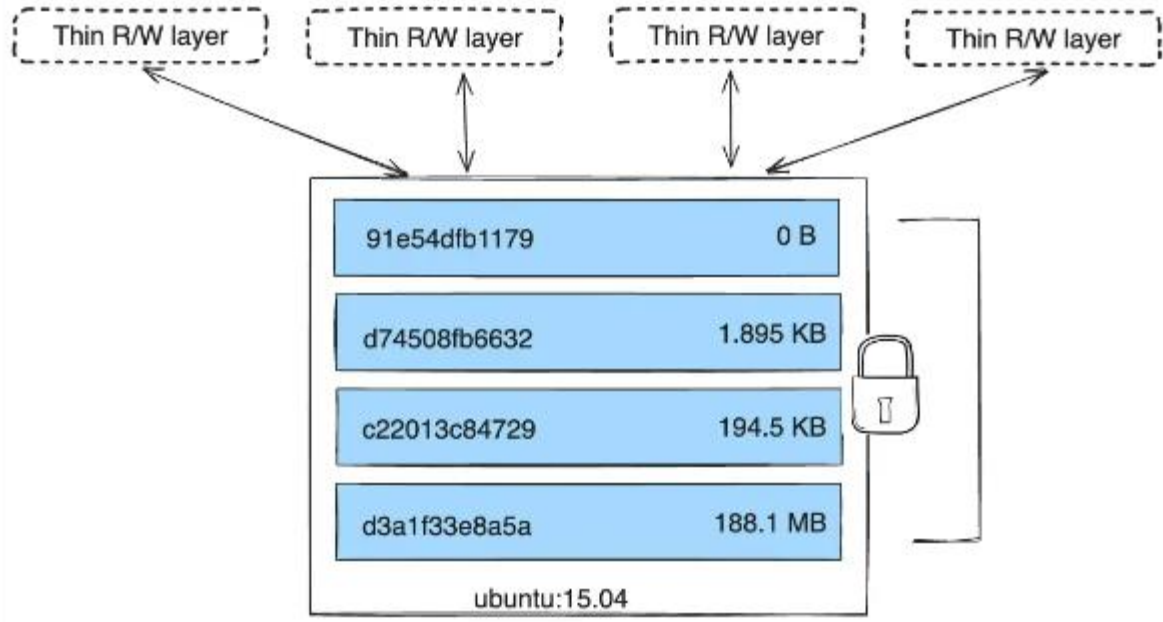
U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	 <p>The diagram illustrates a Docker storage architecture. At the top, four dashed boxes labeled 'Thin R/W layer' are shown. Below them is a central stack of four blue boxes representing image layers. The layers are labeled with their IDs and sizes: '91e54dfb1179' (0 B), 'd74508fb6632' (1.895 KB), 'c22013c84729' (194.5 KB), and 'd3a1f33e8a5a' (188.1 MB). The bottom of the stack is labeled 'ubuntu:15.04'. A padlock icon is positioned to the right of the stack, indicating that the base image is read-only. Arrows show the relationship between the thin R/W layers and the underlying image stack.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>See also description below of containers. See https://docs.docker.com/engine/docker-overview/.</p> <p>On information and belief, a Docker container is a runnable instance of an image and includes configuration options when it is created or started. When an image is run and a container is generated, a writable layer is created (on top of other layer(s)) and includes all changes to the running container.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Containers</p> <p>A container is a runnable instance of an image. You can create, start, stop, move, or delete a container using the Docker API or CLI. You can connect a container to one or more networks, attach storage to it, or even create a new image based on its current state.</p> <p>By default, a container is relatively well isolated from other containers and its host machine. You can control how isolated a container's network, storage, or other underlying subsystems are from other containers or from the host machine.</p> <p>A container is defined by its image as well as any configuration options you provide to it when you create or start it. When a container is removed, any changes to its state that aren't stored in persistent storage disappear.</p> <p>Source: https://docs.docker.com/get-started/docker-overview/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Docker uses storage drivers to store image layers, and to store data in the writable layer of a container. The container's writable layer doesn't persist after the container is deleted, but is suitable for storing ephemeral data that is generated at runtime. Storage drivers are optimized for space efficiency, but (depending on the storage driver) write speeds are lower than native file system performance, especially for storage drivers that use a copy-on-write filesystem. Write-intensive applications, such as database storage, are impacted by a performance overhead, particularly if pre-existing data exists in the read-only layer.</p> <p>On information and belief, a Docker image and a Docker container can be layered in a manner as shown in the image below.</p>  <p>Source: https://docs.docker.com/engine/storage/drivers/overlayfs-driver/.</p>
[7.b.i] said leaf images including only additional data blocks not previously contained in said root image	On information and belief, the Southwest Count 1 Systems and Services practice storing leaf images, where said leaf images including only additional data blocks not previously contained in said root image and changes made by respective compute nodes to the blocks of the root image.

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
and changes made by respective compute nodes to the blocks of the root image,	<p>For example, on information and belief, Docker includes a writable per runnable container layer for new data such as new files and including a copy on write strategy for changes made to the blocks of the root image.</p> <p>The layers are stacked on top of each other. When you create a new container, you add a new writable layer on top of the underlying layers. This layer is often called the "container layer". All changes made to the running container, such as writing new files, modifying existing files, and deleting files, are written to this thin writable container layer. The diagram below shows a container based on an ubuntu:15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<h2>Container and layers</h2> <p>The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted, the writable layer is also deleted. The underlying image remains unchanged.</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	 <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>Docker uses storage drivers to manage the contents of the image layers and the writable container layer. Each storage driver handles the implementation differently, but all drivers use stackable image layers and the copy-on-write (CoW) strategy.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>The copy-on-write (CoW) strategy</p> <p>Copy-on-write is a strategy of sharing and copying files for maximum efficiency. If a file or directory exists in a lower layer within the image, and another layer (including the writable layer) needs read access to it, it just uses the existing file. The first time another layer needs to modify the file (when building the image or running the container), the file is copied into that layer and modified. This minimizes I/O and the size of each of the subsequent layers. These advantages are explained in more depth below.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>On information and belief, when a container is created or started, a thin writable container layer is added on top of the other layers, and any changes the container makes to the filesystem are stored in this layer, whereas any files that the container does not change do not get copied into this layer. Multiple containers can share access to the same image and maintain their own data state using their container layer.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Copying makes containers efficient</p> <p>When you start a container, a thin writable container layer is added on top of the other layers. Any changes the container makes to the filesystem are stored here. Any files the container doesn't change don't get copied to this writable layer. This means that the writable layer is as small as possible.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p><i>See also</i> claim limitations [7.a] and [7.b].</p>

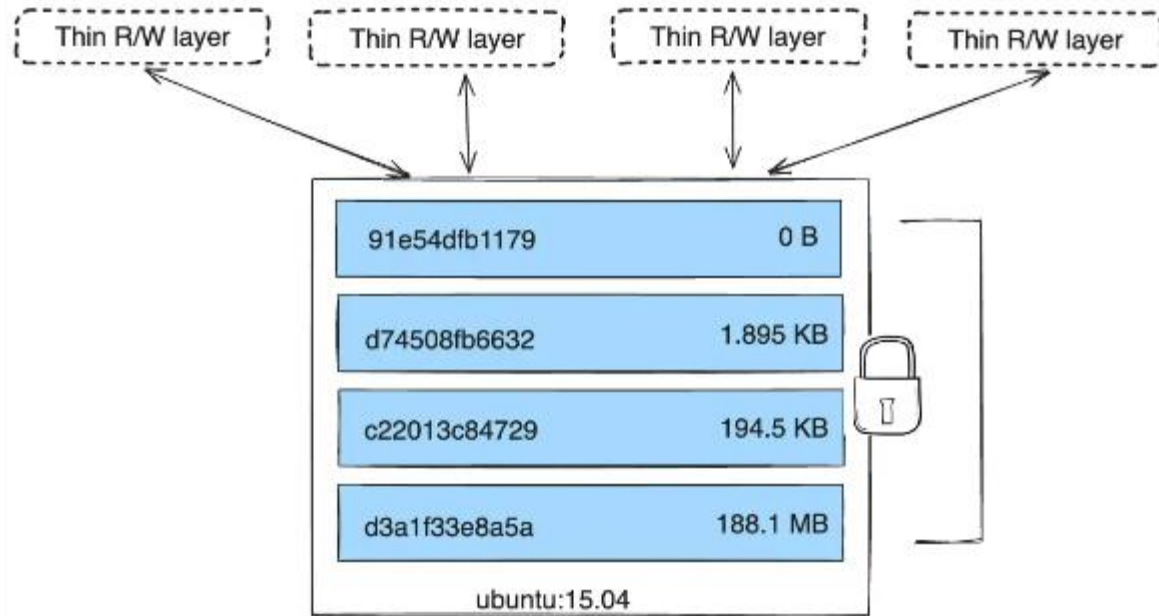
U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
[7.b.ii] wherein said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes; and	<p>On information and belief, the Southwest Count 1 Systems and Services practice said leaf images of respective compute nodes do not include blocks of said root image that are unchanged by respective compute nodes.</p> <p>For example, on information and belief, Docker contains a writable layer using copy on write, leaving unchanged image data in their original image layers.</p> <p>Docker uses storage drivers to manage the contents of the image layers and the writable container layer. Each storage driver handles the implementation differently, but all drivers use stackable image layers and the copy-on-write (CoW) strategy.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <h2>The copy-on-write (CoW) strategy</h2> <p>Copy-on-write is a strategy of sharing and copying files for maximum efficiency. If a file or directory exists in a lower layer within the image, and another layer (including the writable layer) needs read access to it, it just uses the existing file. The first time another layer needs to modify the file (when building the image or running the container), the file is copied into that layer and modified. This minimizes I/O and the size of each of the subsequent layers. These advantages are explained in more depth below.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>On information and belief, when a container is created or started, a thin writable container layer is added on top of the other layers, and any changes the container makes to the filesystem are stored in this layer, whereas any files that the container does not change do not get copied into this layer. Multiple containers can share access to the same image and maintain their own data state using their container layer.</p> <h2>Container and layers</h2> <p>The major difference between a container and an image is the top writable layer. All writes to the container that add new or modify existing data are stored in this writable layer. When the container is deleted, the writable layer is also deleted. The underlying image remains unchanged.</p> <p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p>

U.S. Patent No. 8,332,844 (Claim 7)

Claim(s)

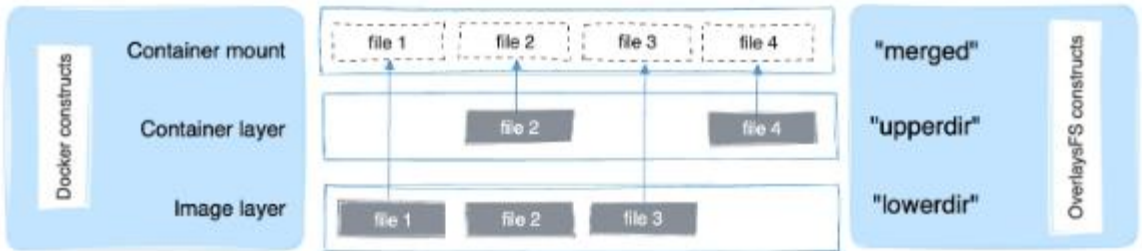
Example Southwest Count 1 Systems and Services

Source: <https://docs.docker.com/engine/storage/drivers/>.

Copying makes containers efficient

When you start a container, a thin writable container layer is added on top of the other layers. Any changes the container makes to the filesystem are stored here. Any files the container doesn't change don't get copied to this writable layer. This means that the writable layer is as small as possible.

Source: <https://docs.docker.com/engine/storage/drivers/>.

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>Because each container has its own writable container layer, and all changes are stored in this container layer, multiple containers can share access to the same underlying image and yet have their own data state. The diagram below shows multiple containers sharing the same Ubuntu 15.04 image.</p> <p>Source: https://docs.docker.com/engine/storage/drivers/.</p> <p>On information and belief, in the example below, the image layer includes file 1, file 2, and file 3, and the container layer includes file 2 and file 4. The container mount provides a unified view of the image and container layers. In this example, the container layer does not include file 1 or file 3, and includes file 2 and file 4, which have either been modified or created.</p>  <p>Source: https://docs.docker.com/engine/storage/drivers/overlayfs-driver/.</p> <p>See also claim limitations [7.a], [7.b], and [7.b.i].</p>
[7.c] caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.	<p>On information and belief, the Southwest Count 1 Systems and Services practice caching blocks of said root image that have been accessed by at least one of said compute nodes in a cache memory.</p> <p>Docker caches unchanged image layers from previous commands.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<h2>Optimize cache usage in builds</h2> <p>When building with Docker, a layer is reused from the build cache if the instruction and the files it depends on hasn't changed since it was previously built. Reusing layers from the cache speeds up the build process because Docker doesn't have to rebuild the layer again.</p> <p>Source: https://docs.docker.com/build/cache/optimize/.</p> <h2>Understanding Docker Cache</h2> <p>Before we dive into the cleaning process, it's important to understand how Docker cache works. When you build a Docker image, Docker uses a build cache to speed up the build process. The build cache stores intermediate layers of the image, which are the layers that don't change frequently. This allows Docker to reuse these layers when building a new image, saving time and resources.</p> <p>Source: https://collabnix.com/how-to-clear-docker-cache/.</p>

U.S. Patent No. 8,332,844 (Claim 7)	
Claim(s)	Example Southwest Count 1 Systems and Services
	<p>At each occurrence of a RUN command in the Dockerfile, Docker will create and commit a new layer to the image, which is just a set of tightly-coupled directories full of various file structure that comprise a Docker image. In a default install, these are located in /var/lib/docker.</p> <p>During a new build, all of these file structures have to be created and written to disk — this is where Docker stores base images. Once created, the container (and subsequent new ones) will be stored in the folder in this same area.</p> <p>What makes the cache important? If the objects on the file system that Docker is about to produce are unchanged between builds, reusing a cache of a previous build on the host is a great time-saver. It makes building a new container really, really fast. None of those file structures have to be created and written to disk this time — the reference to them is sufficient to locate and reuse the previously built structures.</p> <p>Source: https://thenewstack.io/understanding-the-docker-cache-for-faster-builds/.</p>